

Attorney's Docket No. 042390.P6349
Express Mail No. EM560890943US

UNITED STATES PATENT APPLICATION

FOR

82C59 REDIRECTION TO I/O APIC

Inventor:

Rajeev K. Nalawadi

Prepared by:

BLAKELY SOKOLOFF TAYLOR & ZAFMAN LLP
12400 Wilshire Boulevard, Seventh Floor
Los Angeles, California 90025
(310) 207-3800

82C59 REDIRECTION TO I/O APIC

BACKGROUND OF THE INVENTION

5 Field of the Invention

The invention relates to computer system peripheral connections and more particularly to managing system interrupts.

Background

10 Peripheral Component Interconnect (PCI)-Extended Industry Standard Architecture (EISA) bridge sets provide an I/O subsystem for many computer systems. One Peripheral Component Interconnect (PCI) standard is PCI Local Bus standard version 2.2 (January 5, 1999). The chip set generally consists of two
15 components--the PCI-EISA Bridge (PCEB) and the EISA system component (ESC). In general, the ESC implements system functions such as timer-counter, direct memory access (DMA), and interrupt control.

20 In one form of interrupt control utilized in certain chip set configurations of Intel Corporation of Santa Clara, California, an EISA compatible interrupt controller of the ESC incorporates the functionality of two 82C59 interrupt controllers that are cascaded providing fourteen external and two internal interrupts. The ESC also contains an Advanced

Programmable Interrupt Controller (APIC). The APIC can be used in either a uni-processor or multi-processor system. The APIC provides multi-processor interrupt management and incorporates both static and dynamic symmetric interrupt distribution across all processors. In systems with multiple I/O subsystems, each system can have its own set of interrupts.

As noted, the EISA compatible interrupt controller incorporates the functionality of two 82C59 interrupt controllers. The two controllers are cascaded into a master interrupter controller and a slave interrupt controller. Two internal interrupts are used for internal function only and are not available at the chip periphery. One interrupt is used to cascade the two controllers together and another is used as a system timer interrupt. The remaining 14 interrupt lines are available for external system interrupts. The interrupts are programmed to utilize on the order of 2000 logic gates. Examples of uses of the 14 system interrupts include interrupts for a keyboard, hard drive, modem, etc.

While the standard EISA compatible interrupt controller is intended for use in a uni-processor system, the APIC can be used in either a uni-processor or multi-processor system. An APIC provides multi-processor interrupt management and incorporates static and dynamic symmetric interrupt distribution across all

processors. In systems with multiple I/O subsystems, each subsystem can have its own set of interrupts.

At the system level, an APIC consists of two parts--one residing in the I/O subsystem (I/O APIC) and the other in the CPU (local APIC). The ESC contains the I/O APIC unit.

The I/O APIC unit consists of a set of interrupt input signals, a 16-entry Interrupt Redirection Table, programmable registers, and a message unit for sending and receiving APIC messages over the APIC bus. I/O devices inject interrupts into the system by asserting one of the interrupt lines to the I/O APIC. The I/O APIC selects a corresponding entry in the redirection table and uses the information of that entry to format an interrupt request message. Each entry in the redirection table can be individually programmed to indicate edge/level sensitive interrupt signals, the interrupt vector and priority, the destination processor, and how the processor is selected (e.g., statically or dynamically). The information in the table is used to transmit a message to other APIC units via the APIC bus.

In addition to its compatibility with multi-processor environments, the APIC system offers the ability to handle a greater number of system interrupts with greater flexibility than the EISA compatible interrupt controller. Nevertheless, the EISA compatible interrupt controller remains a legacy

[illegible]

	1970-71	1971-72	1972-73	1973-74	1974-75	1975-76	1976-77	1977-78	1978-79	1979-80	1980-81	1981-82	1982-83	1983-84	1984-85	1985-86	1986-87	1987-88	1988-89	1989-90	1990-91	1991-92	1992-93	1993-94	1994-95	1995-96	1996-97	1997-98	1998-99	1999-00	2000-01	2001-02	2002-03	2003-04	2004-05	2005-06	2006-07	2007-08	2008-09	2009-10	2010-11	2011-12	2012-13	2013-14	2014-15	2015-16	2016-17	2017-18	2018-19	2019-20	2020-21	2021-22	2022-23	2023-24	2024-25	2025-26	2026-27	2027-28	2028-29	2029-30	2030-31	2031-32	2032-33	2033-34	2034-35	2035-36	2036-37	2037-38	2038-39	2039-40	2040-41	2041-42	2042-43	2043-44	2044-45	2045-46	2046-47	2047-48	2048-49	2049-50	2050-51	2051-52	2052-53	2053-54	2054-55	2055-56	2056-57	2057-58	2058-59	2059-60	2060-61	2061-62	2062-63	2063-64	2064-65	2065-66	2066-67	2067-68	2068-69	2069-70	2070-71	2071-72	2072-73	2073-74	2074-75	2075-76	2076-77	2077-78	2078-79	2079-80	2080-81	2081-82	2082-83	2083-84	2084-85	2085-86	2086-87	2087-88	2088-89	2089-90	2090-91	2091-92	2092-93	2093-94	2094-95	2095-96	2096-97	2097-98	2098-99	2099-00	2100-01	2101-02	2102-03	2103-04	2104-05	2105-06	2106-07	2107-08	2108-09	2109-10	2110-11	2111-12	2112-13	2113-14	2114-15	2115-16	2116-17	2117-18	2118-19	2119-20	2120-21	2121-22	2122-23	2123-24	2124-25	2125-26	2126-27	2127-28	2128-29	2129-30	2130-31	2131-32	2132-33	2133-34	2134-35	2135-36	2136-37	2137-38	2138-39	2139-40	2140-41	2141-42	2142-43	2143-44	2144-45	2145-46	2146-47	2147-48	2148-49	2149-50	2150-51	2151-52	2152-53	2153-54	2154-55	2155-56	2156-57	2157-58	2158-59	2159-60	2160-61	2161-62	2162-63	2163-64	2164-65	2165-66	2166-67	2167-68	2168-69	2169-70	2170-71	2171-72	2172-73	2173-74	2174-75	2175-76	2176-77	2177-78	2178-79	2179-80	2180-81	2181-82	2182-83	2183-84	2184-85	2185-86	2186-87	2187-88	2188-89	2189-90	2190-91	2191-92	2192-93	2193-94	2194-95	2195-96	2196-97	2197-98	2198-99	2199-00	2200-01	2201-02	2202-03	2203-04	2204-05	2205-06	2206-07	2207-08	2208-09	2209-10	2210-11	2211-12	2212-13	2213-14	2214-15	2215-16	2216-17	2217-18	2218-19	2219-20	2220-21	2221-22	2222-23	2223-24	2224-25	2225-26	2226-27	2227-28	2228-29	2229-30	2230-31	2231-32	2232-33	2233-34	2234-35	2235-36	2236-37	2237-38	2238-39	2239-40	2240-41	2241-42	2242-43	2243-44	2244-45	2245-46	2246-47	2247-48	2248-49	2249-50	2250-51	2251-52	2252-53	2253-54	2254-55	2255-56	2256-57	2257-58	2258-59	2259-60	2260-61	2261-62	2262-63	2263-64	2264-65	2265-66	2266-67	2267-68	2268-69	2269-70	2270-71	2271-72	2272-73	2273-74	2274-75	2275-76	2276-77	2277-78	2278-79	2279-80	2280-81	2281-82	2282-83	2283-84	2284-85	2285-86	2286-87	2287-88	2288-89	2289-90	2290-91	2291-92	2292-93	2293-94	2294-95	2295-96	2296-97	2297-98	2298-99	2299-00	2300-01	2301-02	2302-03	2303-04	2304-05	2305-06	2306-07	2307-08	2308-09	2309-10
--	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------	---------

5 interrupt type is re-routed to a second interrupt controller.

BRIEF DESCRIPTION OF THE DRAWINGS

The features, aspects, and advantages of the invention will become more thoroughly apparent from the following detailed description, appended claims, and accompanying drawings in

5 which:

Figure 1 is a schematic block diagram of a multi-processor environment illustrating the interrupt control for the system.

Figure 2 is a flow chart of an embodiment of the method of the invention.

10 **Figure 3** is a specific embodiment of the method of the invention.

DETAILED DESCRIPTION OF THE INVENTION

In one aspect, the invention relates to a method and system to handle all system interrupts through the APIC system, including processor systems that rely for at least a portion of their interrupt control on the legacy standard of the EISA compatible interrupt controller. In this manner, the invention allows systems that rely partially on EISA compatible interrupt controllers, such as the 82C59 interrupt controller to be incorporated in a multi-processor environment. By utilizing the APIC system of interrupt control the number of dedicated gates required for interrupt control can be reduced.

Figure 1 schematically illustrates a block diagram of a multi-processor environment suitable for interrupt control according to the invention. System 100 is a multi-processor system, including processor 110, processor 210, and processor 310. Each processor contains a memory element (e.g., random access memory, read only memory, cache, etc.) System 100 also includes, in this illustration, EISA system component (ESC) 410. ESC 410 incorporates the standard EISA compatible interrupt controller and the APIC. In one embodiment, the EISA compatible interrupt controller includes two 82C59 interrupt controllers 430 arranged in a master (CNTRL-1)-slave (CNTRL-2) configuration as known in the art. In one embodiment, the two 82C59 interrupt

controllers are utilized to control the interrupt functions listed in Table 1:

Table 1: Typical Interrupt Functions

Priority	Label	Controller	Typical Interrupt Source
1	IRQ0	1	Interval Timer 1, Counter 0 OUT
2	IRQ1	1	Keyboard
3-10	IRQ2	1	Interrupt from Controller 2
3	IRQ8#	2	Real Time Clock
4	IRQ9	2	Expansion Bus Pin B04
5	IRQ10	2	Expansion Bus Pin D03
6	IRQ11	2	Expansion Bus Pin D04
7	IRQ12	2	Expansion Bus Pin D05
8	IRQ13	2	Coprocessor Error, Chaining
9	IRQ14	2	Fixed Disk Drive Controller Expansion Bus Pin D07
10	IRQ15	2	Expansion Bus Pin D06
11	IRQ3	1	Serial Port 2, Expansion Bus B25
12	IRQ4	1	Serial Port 1, Expansion Bus B24
13	IRQ5	1	Parallel Port 2, Expansion Bus B23
14	IRQ6	1	Diskette Controller, Expansion Bus B22
15	IRQ7	1	Parallel Port 1, Expansion Bus B21

The following demonstrates the interrupt sequence for a 80x86-type system.

1. One or more of the Interrupt Request (IRQ[x]) lines are raised high, setting the corresponding Interrupt Request Register (IRR) bit(s).
2. The Interrupt Controller evaluates these requests, and sends an INTR to the CPU, if appropriate.
3. The CPU acknowledges the INTR and responds with an interrupt acknowledge cycle. This cycle is translated

into a PCI bus command. This PCI command is broadcast over the PCI bus as a single cycle.

4. Upon receiving an interrupt acknowledge cycle from the CPU over the PCI, the PCI-EISA bridge (PCEB) converts the single cycle into an INTA# pulse to the ESC. The ESC uses the INTA# pulse to generate the two cycles that the internal 82C59 pair can respond to with the expected interrupt vector. The cycle conversion is performed by a functional block in the ESC Interrupt Controller Unit. The internally generated interrupt acknowledge cycle is completed as soon as possible as the PCI bus is held in wait states until the interrupt vector data is returned. Each cycle appears as an interrupt acknowledge pulse on the INTA# pin of the cascaded interrupt controllers. These two pulses are not observable at the ESC periphery.
5. Upon receiving the first internally generated interrupt acknowledge, the highest priority In-Service Register (ISR) bit is set and the corresponding IRR bit is reset. The Interrupt Controller does not drive the Data Bus during this cycle. On the trailing edge of the first cycle pulse, a slave identification code is broadcast by the master to the slave on a private, internal three bit wide bus. The slave controller uses these bits to

determine if it must respond with an interrupt vector during the second INTA# cycle.

6. Upon receiving the second internally generated interrupt acknowledge, the Interrupt Controller releases an 8 bit pointer (the interrupt vector) onto the Data Bus where it is read by the CPU.

7. This generally completes the interrupt cycle. In an automatic end-of-interrupt mode the ISR bit is reset at the end of the second interrupt acknowledge cycle pulse. Otherwise, the ISR bit remains set until an appropriate end-of-interrupt (EOI) command is issued at the end of the interrupt subroutine.

At the system level, the APIC consists of two parts -- one residing in the I/O subsystem (I/O APIC 420) and the other, local APIC 120, 220, and 320, residing in processor 110, processor 210, and processor 310, respectively. The local APIC and I/O APIC 420 communicate over dedicated APIC bus 400. The ESC's I/O APIC bus interface consists of bi-directional data signals and a clock input. Each local APIC (local APIC 120, local APIC 220, and local APIC 320) contains intelligence to determine whether or not each processor should accept interrupts broadcast on APIC bus 400. The local APIC unit also provides local pending of interrupts, nesting and masking of interrupts, and handles all interactions with its local processor (e.g.,

interrupt request (INTR), interrupt acknowledge (INTA, and EOI protocol). Each local APIC also provides inter-processor interrupts and a timer to its local processor.

As shown in **Figure 2**, the invention contemplates identifying and redirecting EISA compatible interrupt signals (block 450), masking the EISA compatible interrupt controller interrupt addresses (block 460) and re-directing the interrupt addresses to the I/O APIC (block 470). In this manner, some or all of the interrupts handled by the EISA compatible interrupt controller (e.g., the 82C59 controller(s)) are instead handled by the I/O APIC.

Figure 3 describes a method of redirecting 82C59 interrupt addresses to an I/O APIC. This embodiment takes advantage of the normal process of initializing interrupt registers on system start-up. The embodiment seeks to identify 82C59 interrupt addresses through system logic. One way this is accomplished is by configuring a system management interrupt (SMI) to trap a configuring 82C59 address (block 510). An SMI is derived from power management architecture of the ESC. In particular, the system management mode (SMM) is a function of the ESC's power management architecture that includes software that controls transitions between a Power-On state and a Fast-Off state. The Fast-Off state is generally utilized in modern processors to permit a system to operate in a low power state without being

powered down. When the system is in a Fast-Off state, the system consumes less power than the Power-On state and is not in a Power Off state. Events such as pressing a key on a keyboard or moving a mouse constitutes a Fast-On event to convert from a Fast-Off state to a Power-On state.

To invoke SMM software, the ESC generates an SMI to the processor. An SMI offers serialized execution. Once an SMI happens, the processor stops and processes the SMI.

The 82C59 generally contains four initializing address registers -- 20h, 21h, A0h, A1h. The four address registers are initialized by four initialization command words: ICW1, ICW2, ICW3, and ICW4. Before normal operation can begin, each interrupt controller in the system must be initialized. In the ESC, this is a four byte sequence. The base address for each interrupt controller is a fixed location in I/O memory space, at 0020h for CNTRL-1 and 00A0h for CNTRL-2. An I/O write to the CNTRL-1 or CNTRL-2 base address with data bit 4 equal to 1 is interpreted as ICW1. For ESC-based EISA systems, three I/O writes to "base address + 1" (021h for CNTRL-1 and 0A0h for CNTRL-2) must follow the ICW1. The first write to "base address + 1" (021h/0A0h) performs ICW2, the second write performs ICW3, and the third write performs ICW4. Initialization of the 82C59 controller thus generally occurs in a particular sequence, ICW1, ICW2, ICW3, and ICW4.

In one embodiment, an SMI is configured to recognize and trap a configuring 82C59 address, such as one of the identified initialization command words. In addition to trapping a configuring 82C59 address, the trapped address is analyzed to determine whether or not the address is the first configuring address (block 520). As noted above, 82C59 is configured in a specific sequence, therefore, prior to re-routing the initialization addresses to I/O APIC as a trap, the SMI locates the first configuring address. A look-up table containing the appropriate order of the initialization command words is provided to identify the appropriate configuring address.

Once the SMI locates the first configuring address, the SMI then locates the remaining configuring addresses (block 530). An SMI vector handler is then generated to direct the appropriate processor to find the interrupt (block 540). The SMI vector handler is associated with the I/O APIC. Once the SMI vector handler is generated, an SMI vector is passed to the appropriate processor (block 550). In response to the SMI vector, the processor re-routes configuring addresses to the I/O APIC (block 560). In this regard, the I/O APIC is then configured to handle 82C59 interrupts.

By handling all 82C59 interrupts through the APIC system, the invention supports the use of conventional chip set systems with advanced operating systems that do not support 82C59

[illegible]

5